

TITLE OF THE INVENTION

Efficient Revocation of Registration Authorities

5

CROSS REFERENCE TO RELATED APPLICATIONS

N/A

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR

10 DEVELOPMENT

N/A

BACKGROUND OF THE INVENTION

The present invention relates to security mechanisms
15 within computer networks and more specifically, to a method and system for efficiently revoking a certificate issued by a certification authority upon a request from a registration authority.

The use of certification authorities and registration authorities within computer networks is well known. A certification authority issues certificates that may be relied upon by other parties that trust the respective certification authority. For example, a certification authority may issue an identity certificate that includes the identity of a principal and a public key associated with the principal. The certification authority issuing the certificate typically authenticates the certificate. Alternatively, a certificate may comprise a group membership certificate that identifies members of a group, the identity of subgroups that

identify members of the group, or non-members of the group.

In a number of circumstances a hierarchical structure is established in which information used by a certification authority to grant a certificate is obtained from one or more registration authorities (RAs).
5 Multiple RAs may be provided in a system for the convenience of the principals. For example, it may be desirable to have an RA at each corporate office but only
10 have certificates issued by a single CA. In this circumstance information necessary for the CA to issue a certificate must be obtained from the respective RA.

The use of RAs in association with CAs for the generation and issuance of certificates is well known in
15 the art. Such systems often employ the Public Key Infrastructure (PKI) and rely on the use of public/private key pairs for purposes of authentication. Only the certification authority may know the RA public keys. A principal can request a certificate from one of
20 a plurality of RAs that is accessible by the principal. The RA, in response to the request from the principal, forwards a request to the CA to issue a certificate for the principal. The request may be digitally signed using the private key of the respective RA. Upon verification
25 by the CA that the request is authentic, the CA issues the certificate. One known Certificate Request Message Format (CRMF) is described in the Internet X.509 Certificate Request Message Format specification dated March 1999 and described in Request for Comments (RFC)
30 2511.

Subsequent to the issuance of a certificate by a CA based upon information provided by an RA, however, it may be determined that the RA has become untrustworthy. While it is straightforward to prevent the CA from issuing further certificates based upon information provided by the untrustworthy RA, it is not easy to revoke previously issued certificates based upon information provided by the untrustworthy RA. The CA may revoke the certificates issued by the untrustworthy RA.

5 The certificates, however, must be revoked individually and each certificate must be listed in and tested against a potentially large certificate revocation list (CRL) to ascertain whether the specific certificate is contained on the respective list. This can be a time consuming

10 process that adds latency to the determination of whether

15 a particular certificate has been revoked.

Various techniques have been proposed for managing the certificate revocation process. Two such techniques are described in U.S. Patents 5,261,002 and 5,687,235.

20 Such techniques, however, do not address the problem of how to efficiently revoke certificates upon recognition that a particular RA has been untrustworthy.

It would therefore be desirable to have an efficient mechanism for revoking certificates issued by a CA at the request of an RA that has been determined to be untrustworthy.

BRIEF SUMMARY OF THE INVENTION

Consistent with the present invention, a method and

30 system for efficiently revoking certificates that were

generated by a certification authority (CA) in response to a request from a registration authority is disclosed. Upon receipt of sufficiently trustworthy information from or on behalf of a principal that requests issuance of a certificate for the respective principal, the registration authority generates a certificate request message (CRM) on behalf of the principal and forwards the CRM to a certification authority. The CRM typically includes the identity of the principal and the identity of the RA and may be authenticated by the RA that generated the request. The CA, upon receipt of the CRM from the RA, in a preferred embodiment, generates a certificate that includes the identity of the principal. The principal may comprise an individual, a client, a server, a software process, identifiable hardware or a system component, or a group. The certificate also includes an RA identifier associated with the RA that forwarded the respective CRM. Optionally, the certificate may include the time at which the CRM was forwarded by the respective CRM to the CA.

In response to a determination that the RA that requested issuance of the respective certificate has become untrustworthy, the CA may generate an entry within a Certificate Revocation List (CRL) in the form of an RA identifier that identifies the untrustworthy RA.

In response to a request for service or access to a resource received at a server from a principal, a determination is made whether the principal is authorized to obtain the requested service. During the authentication process, the server accesses a certificate

me
10/21/03

associated with the principal that includes a public key
~~key~~ associated with the principal. The public key may be
used by the server to verify the principal's request.
Additionally, a determination is made whether the RA
5 identifier contained within the respective certificate
matches an RA identifier on the CRL prior to granting
access to the requested service or resource. In the
event the RA identifier within the certificate matches
the RA identifier within the CRL, an indication is
10 provided to the server that the certificate has been
revoked. In response to this indication, the server may
deny service to the requesting principal. Alternatively,
if the determination reveals that the RA identifier
contained within the respective certificate is not
15 contained on the CRL, the server may grant access to the
requested service or resource or perform additional
validations pertaining to the request prior to granting
such access.

In the event that the RA has become untrustworthy,
20 in addition to the RA identifier that is added to the
CRL, a date or dates may be included in the CRL in
association with the RA identifier. The date(s) specify
a period or period(s) for which certificates issued by a
CA at the request of the respective RA are deemed to be
25 untrustworthy. More specifically, certificates issued
within such periods are deemed to be revoked. For
example, a single date may be associated with an RA
identifier in the CRL. All certificates requested by the
respective RA after the date specified within the CRL may
30 be considered to be revoked. Additionally, multiple

dates defining beginning and ending times of a period or periods may be employed to identify certificates that have been issued by the CA in response to CRMs from an RA at times when the RA is deemed to have been untrustworthy and thus revoked.

Other features, aspects and advantages of the presently disclosed method and system will be apparent from the Detailed Description of the Invention that follows.

10

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWING

The invention will be more fully understood by reference to the Detailed Description of the Invention in conjunction with the Drawings, of which:

15

Fig. 1 is a block diagram illustrating a system operative in a manner consistent with the present invention;

20

Fig. 2 is a block diagram of an exemplary computer system that may be employed to perform the functions of the client, the servers, registration authorities and certification authorities depicted in Fig. 1;

25

Fig. 3 is a flow diagram illustrating an exemplary method of operation of the system depicted in Fig. 1 for generating a certificate in a manner consistent with the present invention;

Fig. 4 is a flow diagram illustrating an exemplary method of operation of the system depicted in Fig. 1 for revoking a certificate;

30

Fig. 5 is a diagram of an exemplary certificate request message forwarded from a registration authority

to a certification authority to request issuance of a certificate for a principal; and

Fig. 6 is a diagram illustrating a certificate that includes a registration authority identifier and a time stamp that identifies the time associated with the certificate request message.

DETAILED DESCRIPTION OF THE INVENTION

A system 10 for efficiently revoking certificates issued by a certification authority upon a request from a registration authority is depicted in Fig. 1. The system 10 includes a plurality of clients 12 identified as clients 12_a through 12_n, at least one service such as provided by a file server 14, a plurality of registration authorities (RAs) 16 identified as RAs 16_a through 16_n and at least one certification authority (CA) 18. A plurality of certification authorities designated 18_a through 18_n are illustrated. Additionally, the system 10 includes a directory server (DS) 20 that is operative to provide access to certificates issued by one or more of the CAs 18 and a revocation server (RS) 22 that maintains one or more certificate revocation lists (CRLs). The clients 12, the file server 14, the RAs 16, the CAs 18, the directory server 20 and the revocation server 22 are communicably coupled via a network 24 to allow the communication of information and/or messages between the respective devices. The network 24 may comprise a local area network, a wide area network, a global computer network such as the Internet, or any other network for communicatively coupling the respective devices.

The clients 12, the file server 14, the RAs 16, the CAs 18, the directory server 20 and the revocation server 22 each typically comprise a computer system 30 such as depicted generally in Fig. 2. The computer system 30 may 5 be in the form of a personal computer or workstation, a personal digital assistant (PDA), an intelligent networked appliance, a controller or any other device capable of performing the functions attributable to the respective devices as described herein.

More specifically, referring to Fig. 2, the computer 10 system 30 typically includes a processor 30a that is operative to execute programmed instructions out of an instruction memory 30b. The instructions executed in performing the functions herein described may comprise 15 instructions stored within program code considered part of an operating systems 30e, instructions stored within program code considered part of an application 30f or instructions stored within program code allocated between the operating system 30e and the application 30f. The 20 memory 30b may comprise random access memory or a combination of random access memory and read only memory. Each device within the system 10 includes a network interface 30d for coupling the respective device to the network 24. The devices within the system 10 may 25 optionally include secondary storage 30c.

The operation of the system 10 may be considered in two phases. The first phase is illustrated in the flow diagram of Fig. 3 and involves the generation of a certificate on behalf of a principal and the second 30 phase, illustrated in the flow diagram of Fig. 4,

involves the use of the certificate in a determination of whether access to a resource or service accessible via the network 24 should be made available to the requesting principal. As described above, for purposes of the present discussion, the term "principal" is intended to refer to nodes within the computer network such as a client or a server, a software process running on a network node, a user or any other component within the network that is capable of requesting access to a service or resource available via the network 24.

More specifically, referring to Fig. 3, an RA 16 receives a request for issuance of a certificate on behalf of a principal as depicted in step 70. For purposes of illustration, it is assumed that the principal comprises client_a 12a and the RA 16 comprises RA_a 16_a. The principal may provide the request directly to RA_a 16_a, or alternatively, the request may be provided to RA_a 16_a by a system administrator (not shown). In response to the request to RA_a 16_a for issuance of a certificate for the principal, as illustrated in step 72, RA_a 16_a may make a determination whether a certificate should be issued for the principal, namely client 12_a. The determination may comprise an analysis of credentials accompanying the request, verifying the authenticity of the request, or any other suitable basis for determining whether the certificate should be issued for the principal. In the event it is determined in inquiry step 72 that no certificate should be issued, no certificate is generated and the process of certificate generation terminates as illustrated in step 82. In the event it is

determined in inquiry step 72 that a certificate should be issued for the principal, a certificate request message (CRM) is forwarded from RA_a 16_a to a certification authority 18. For purposes of the present example, it is 5 assumed that the CRM is forwarded from registration authority RA 16_a to CA_a 18_a.

An illustrative CRM 40 is depicted in Fig. 5. Referring to Fig. 5, the CRM 40 typically includes at least a certificate request portion 42 and an 10 authentication portion 44. The certificate request portion 42 comprises a request from the respective RA 16 to the respective CA 18 that a certificate be issued for the principal identified in the certificate request portion 42 e.g. client 12_a in the present example. When 15 the public key infrastructure (PKI) is being employed, the authentication portion 44 may comprise a digital signature in which the certificate request message 40 is signed by the RA 16 using the RA's respective private key.

Referring again to Fig. 3, upon receipt of the CRM 40 at CA_a 18, a determination is made by CA_a 18_a whether the request received from RA_a 16_a is a valid request as depicted in decision step 76. More specifically, if CA_a 18_a determines that the request is not a valid request, 20 CA_a 18_a does not generate a certificate for the respective principal and certificate generation terminates as depicted in step 82. In the event that CA_a 18_a determines that the CRM comprises a valid request, certificate 25 generation continues, as depicted in step 78. The process of verifying the CRM may comprise the step of 30

verifying the authenticity of the CRM 40 by using the public key of RA_a 16_a to check a digital signature included in the CRM 40. Alternatively, any other suitable technique for authenticating the CRM 40 may be 5 employed. Additionally, the CA_a 18_a may optionally verify other credentials pertaining to the CRM or the principal or perform other tests prior to generation of a certificate for the principal.

Upon determining that the CRM 40 comprises a valid 10 request for issuance of a certificate, CA_a 18_a generates the certificate as depicted in step 78. An exemplary certificate issued by a CA 18, such as CA_a 18_a, in response to a valid CRM 40 is illustrated in Fig. 6.

Referring to Fig. 6, the certificate 50 includes a 15 principal identifier 52 associated with the respective principal, a principal public key 54 associated with the principal identifier 52, and an RA identifier 56 that identifies the respective RA 16 that forwarded the CRM 40 to the respective CA 18 requesting issuance of the 20 certificate 50. Additionally, the certificate 50 may optionally include a time stamp 58 that indicates the time when the CRM 40 was received by the CA 18. The certificate 50 further includes an authentication portion 25 60 that may comprise the digital signature of the CA 18 issuing the certificate 50 or any other suitable form of authentication. By way of illustration, it is assumed that the certificate 50 includes a principal identifier for client_a 12_a, the public key associated with a private key owned by client_a 12_a, an RA identifier for RA_a 16_a and 30 a time stamp that specifies the time when the respective

CRM 40 was received by CA_a 18_a from RA_a 16_a. In the illustrative example, the certificate 50 is digitally signed by CA_a 18_a using the private key owned by that CA.

5 The certificate 50 generated in the above-described manner is published by CA_a 18_a as illustrated in step 80. Publication may involve transmittal of the certificate 50 to a directory server 20 (Fig. 1) that maintains certificates 50 generated by CA_a 18_a. Alternatively, the certificates may be delivered to the respective
10 principal. Any other suitable technique known in the art for publishing or distributing the certificates 50 may also be employed.

15 The use of the certificate 50 during system 10 operation is described below with respect to Figs. 1, 4 and 6.

20 When a principal desires to obtain access to a service or resource accessible via the network 24, the principal generates a request for the identified service or resource and transmits the request over the network 24 to the applicable server as depicted in step 90. For purposes of illustration, it is assumed that client_a 12_a desires to access a file maintained on the file server 14 (Fig. 1). While the resource is depicted to be a file stored on a file server 14, it should be appreciated that
25 client_a 12_a or any other principal may be attempting to obtain access to any service or resource accessible via the network 24. Upon receipt of the request at the file server 14, the file server 14 obtains the certificate 50 for the principal, e.g. client_a 12_a as depicted in step
30 92. The certificate 50 may be stored locally, may be

obtained from the principal, or may be obtained from the directory server 20. The file server 14 determines whether the request received from the respective principal comprises a valid request as depicted in
5 inquiry step 94. For example, the file server 14 may obtain the certificate 50 for client_a 12_a from the directory server 20 and utilize a public key associated with client_a 12_a that is contained within the respective certificate to verify a digital signature in the request
10 from client_a 12_a. In the event the file server 14 determines that the request from client_a 12_a is not a valid request, the file server 14 denies access to the requested file as illustrated in step 96. In the event the file server 14 determines that the request is a valid
15 request, the file server 14 determines whether the certificate for the respective principal has been revoked. In this regard, the file server 14 accesses a copy of the CRL as depicted in step 98. More specifically, the file server 14 retrieves a recent copy
20 of the CRL if it does not possess a recent copy via any suitable CRL distribution or publishing technique known in the art. The distribution of CRLs identifying untrustworthy RAs throughout the network has several advantages. First, since many certificates may be issued
25 at the request of a single RA, all of the certificates that are deemed untrustworthy need not be separately identified in the CRL. Rather, such certificates may be identified via a single entry in the CRL. Accordingly, processing resources needed to update the CRL are
30 reduced. Additionally, the CRL is smaller in size, less

bandwidth is required to distribute the CRL to various services throughout the network and the CRL occupies less storage space once distributed to each of the various services.

5 A determination is made, as depicted in step 99, whether the certificate has been explicitly revoked. In the event the certificate has been explicitly revoked, control passes to step 104. If the certificate has not been explicitly revoked, control passes to step 100.

10 In the event it is determined that either the certificate has been explicitly revoked, as determined in step 99, or that the certificate has been issued by an RA that has been deemed untrustworthy, as indicated per steps 100 and 102, access to the requested service may be denied as illustrated in step 104.

15 In an alternative embodiment, the file server 14 may forward the certificate to the revocation server 22 and the revocation server may analyze a CRL maintained at the revocation server 22 to determine whether the certificate 20 50 has been explicitly revoked as depicted in step 99 or if the RA identified in the certificate matches an RA identifier on the CRL as depicted in step 100. In such event, the revocation server 22 provides an indication to the file server 14 indicating whether the certificate was 25 explicitly revoked or whether the certificate contained an identifier of an RA that has been deemed untrustworthy.

As indicated above, the CRL may include RA identifiers of RAs 16 that have been deemed to be untrustworthy. An entry on the CRL identifying an

untrustworthy RA may optionally include a time threshold that specifies when the respective RA became untrustworthy. The entry may be readily employed to revoke certificates issued by a CA 18 in response to a
5 request by the respective RA 16 as described below.

Assume in a first example, that RA_a 16_a has become untrustworthy, and that it is desired to revoke all certificates that were issued by RA_a 16_a. An entry in the CRL identifying an RA identifier for RA_a 16_a would be
10 inserted in the CRL. In response to an inquiry from the file server 14 a determination is made whether the RA identifier contained within the respective certificate is contained within the CRL as illustrated in step 100. In response to a determination that the RA identifier
15 contained within the certificate 50 is not identified on the CRL, an indication is provided that the respective certificate 50 has not been revoked as illustrated in step 108. In response to this indication, the principal, e.g. client_a 12_a is provided access to the requested file.
20 In the event the RA identifier contained within the respective certificate is contained in the CRL (assuming for the present example no time stamp is employed) control passes to step 104. As indicated in step 104, an indication is provided that the certificate has been
25 revoked and, as indicated in step 106, access to the requested resource (file) is denied.

In a further example, it is assumed that it has been determined that at time_x, RA_a 16_a has become untrustworthy and accordingly, it is desired to revoke all certificates
30 issued by RA_a 16_a after time_x. In such event, an entry

may be made in the CRL that identifies the respective RA that became untrustworthy along with an indication of the time when the respective RA became untrustworthy. For example, assuming RA_a 16_a became untrustworthy as of time_x, an entry on the CRL may be provided as follows:

RA_a time_x

Thus, returning to Fig. 4, the CRL is analyzed to determine whether the RA identifier contained within the certificate 50 corresponds to an entry in the CRL as depicted in step 100. In the event the RA identifier contained within the respective certificate 50 is listed on the CRL, control passes to inquiry step 102. As indicated in inquiry step 102, a determination is made whether the time stamp within the certificate 50 that indicates when RA_a 16_a requested issuance of the respective certificate is after time_x contained within the respective entry on the CRL. In the event the certificate 50 was requested by RA_a 16_a after time_x when the respective RA 16 was determined to have become untrustworthy, an indication is provided to the file server 14 that the certificate 50 is untrustworthy or has been revoked as illustrated in step 104, and access to the requested resource or service is denied as depicted in step 106. In the event the time stamp contained within the respective certificate 50 was generated before time_x, an indication is provided to the file server 14 that the certificate 50 has not been revoked, as indicated in step 108, and access to the requested resource is provided if the principal satisfies other

applicable access control requirements as illustrated in step 110.

In a further example, the CRL may contain a number of time periods in which the respective RA has been determined to be untrustworthy, and an indication may be provided to the file server 14 that a certificate 50 has been revoked if the respective certificate 50 was requested by the untrustworthy RA 16 during any period in which that RA has been determined to be untrustworthy.
5
For example, an entry in the CRL may be provided as follows:
10

RA_a time_{s1}, time_{e1}, time_{s2}, time_{e2}, time_{s3}
where time_{s1} indicates the beginning of the first period in which RA_a 16_a was determined to have become untrustworthy, time_{e1} indicates the end of the first period in which RA_a 16_a was determined to have become untrustworthy, time_{s2} indicates the beginning of the second period in which RA_a 16_a was determined to have become untrustworthy, time_{e2} indicates the end of the second period in which RA_a 16_a was determined to have become untrustworthy, and time_{s3} indicates the beginning of the third period in which RA_a 16_a was determined to have become untrustworthy. Since no end time is provided for the third period, any certificates 50 requested by RA_a 16_a after time_{s3} are considered untrustworthy. More specifically, in the event a single starting time stamp is provided, the end of the period in which the respective RA 16 is considered to be untrustworthy is assumed to be the present time.
15
20
25

In the foregoing manner, an efficient mechanism for revoking certificates issued by a CA in response to a request from an RA is provided. The presently described technique permits CRLs to be generated more quickly and 5 produces smaller CRLs. The smaller CRLs utilize less bandwidth during distribution and require less memory to store. Such memory savings are magnified by the number of services that store a copy of the CRL. Accordingly, overall performance of the system is improved.

10 Those skilled in the art should readily appreciate that the programs defining the functions performed by the respective devices described herein can be communicated to the respective devices in many forms including, but not limited to: (a) information permanently stored on 15 non-writable storage media (e.g., read-only memory devices within a computer such as ROM or CD-ROM disks readable by a computer I/O attachment; (b) information alterably stored on writable storage media (e.g., floppy disks, tapes, read/write optical media and hard drives); 20 or (c) information conveyed to a computer through a communication media, for example, using baseband signaling or broadband signaling techniques, such as over computer or telephone networks via a modem. In addition, while in the present embodiment the functions are 25 illustrated as being software-driven and executable out of a memory by a processor, the presently described functions may alternatively be embodied in part or in whole using hardware components such as Application Specific Integrated Circuits (ASICs), programmable logic arrays, state machines, controllers or other hardware 30

components or devices, or a combination of hardware components and software.

While the above-described examples illustrate a technique for accessing a file on a file server and the
5 use of a certificate including an RA identifier in that process, it should be appreciated that the presently disclosed methods and systems may be used for determining whether access should be provided to any suitable service or resource accessible over a network such as a web page,
10 a secure area, data within a database or privileges within a computer network.

Additionally, it should be appreciated that the authentication techniques described hereinabove may involve digital signatures based upon public/private key
15 pairs as employed within the public key infrastructure (PKI), other asymmetric key pairs or symmetric keys. Additionally, authentication may be performed using a keyed hash, any suitable cryptographic hash incorporated in an encrypted message or any other suitable
20 authentication technique known in the art.

Moreover, while the term certificate, as used herein, is intended to include traditional certificates, such as identity or group certificates that include an identifier of a party or group and an associated public
25 key, the term certificate, is also intended to encompass any document or data structure that is issued at the request of a first party by a second party and that contains an identifier indicative of the identification of the first party, whether or not the certificate is
30 authenticated by the second party. By way of example and

not limitation, a certificate may include an identifier for a party and the name of group a group in which the party is a member. Additionally, a certificate may include the name of a party and a dollar amount that the
5 party is authorized to sign for.

Finally, it will be appreciated by those of ordinary skill in the art that modifications to and variations of the above-described methods and system for efficiently revoking revoking certificates generated at the request
10 of a first node by a second node may be made without departing from the inventive concepts described herein. Accordingly, the invention should not be viewed as limited except as by the scope and spirit of the appended claims.
15